Remote Administration

David M. Ihnat
dihnat@dminet.com

December, 2003

1. INTRODUCTION

This memorandum discusses the issues surrounding implementation of full remote administration capabilities for a typical small IT installation, including a high-level overview of requirements and issues surrounding remote administration.

2. OVERVIEW

Remote administration refers to the ability to provide the following capabilities to an IT facility:

- **Security: Authenticated Remote Access**
- **Remote Notification**.

The capability to provide full remote administration of computer networks, servers, and support systems and workstations provides a number of advantages; some of these advantages—although certainly not an exhaustive list—are:

- **Prompt Notification and Response**. Even when personnel are off-duty or off-site, notification
- **Wider Coverage with Less Inconvenience**. Administrative personnel can provide fast response to a great range of issues from virtually any location. That "call in the night" may only take a few minutes from an Internet-connected computer, instead of a drive to and from the site.
- **Outsourcing**. With the ability to provide audited, authenticated access to administrative functions, the possibility exists to off-load work to outsourced contractors or support organizations, or to even outsource administration of the entire site.

3. REMOTE ADMINISTRATION

3.1 SECURITY

There is no question that the access method-of-choice today is via an Internet TCP/IP connection. Internet access is ubiquitous, fast, and widely supported.

There is also no question that the top concern today when considering an Internet connection over which privileged operations—and system administration heads that list--is security. How do you allow connections from virtually anywhere in the world, over a common communications cloud like the Internet, and provide acceptable assurance that:

Ø The connection is only established with authorized individuals

Copyright ® 2003, 2004 David M. Ihnat

- Ø The connection is secure
- Ø The user is, in fact, an authorized individual

These issues are addressed in the following sections.

3.1.1 Secure Remote Access

Even before an individual is to be permitted to attempt access to the servers and components of a network, a link must be established. In the days when modems were the best (or only) option, the only real concern was whether to allow the connection—was the calling entity actually entitled to establish the link? Once that was established[1], a private point-to-point circuit was assumed; the actual information passing over the dialup connection was considered secure.[2] In fact, if modem access is acceptable[3], this model still applies today.

Typically, however, high-speed access is required, particularly for the full-screen GUI interface required for such systems as Windows servers.

3.1.1.1 Virtual Private Networks

The solution settled upon by the industry is the Virtual Private Network (VPN). When utilizing a VPN, the contents of the virtual circuit established between the requestor and the site is encrypted. Even if someone were to intercept the individual packets that comprise the session, it would do them no good—the *payload*, that part of the packet that contains the actual information being exchanged, is unintelligible.

3.1.1.1.1 Router/Firewall VPN

A VPN may be established in several ways. One of the most common approaches requires the network equipment responsible for fielding connection requests from the Internet—usually the corporate network router—to be capable of establishing a VPN[4]. Once such a connection is established, it appears to the remote system that it is directly connected to the internal LAN. From that point, anything that can be accomplished by a workstation on the LAN can be carried out on the remotely connected workstation.

These connections can configured in one of three ways:

- Between router/firewalls from the same vendor on both ends, in which case the remote user effectively does nothing special to make use of the connection
- Between router/firewalls from different vendors that “know” how to cooperate

[1] Usually via a password required by the modem, e.g., the US Robotics Courier V.Everything modem

[2] Of course, techniques such as wiretapping could compromise modem communications. Largely, however, this has never been considered a meaningful threat in any but the most high-security installations.

[3] Usually, speed is a limiting factor, although modem speeds may be tolerable as an emergency backup connection in case a high-speed Internet link isn’t available.

[4] Cisco, for instance, provides a proprietary VPN solution; the router can be programmed on the corporate end, and Cisco provides client software that the remote system attempting the connection must have loaded. Typically, if the router is provisioned for VPN service, these clients are provided gratis, or for very nominal fees.

- Between a router/firewall and client software (usually provided by the hardware vendor, but sometimes—especially in the case of Linux systems—available from other sources) on the remote system or workstation.

3.1.1.1.2 Tunneled VPN

Another approach, when the firewall/router can't be used for some reason, is to permit a *tunneled request* from authorized external IP addresses to be directed by the corporate router to a server for which the operating system supports VPN connections. A Windows VPN connection could be established in this manner, as could one with a Linux system running any one of several VPN solutions (such as CIPE or SSH). Once connected in this manner, the remote system can make use of whatever resources are available on the host server to access the internal LAN, including administrative, communications, and remote control utilities.

Most of the software required to establish encrypted sessions is available as either commercial or Open Source; both have proven to be reliable and viable solutions.

3.1.2 Authentication

There are two levels of authentication when considering remote administration—connection authentication, and service authentication once the connection is established.

3.1.3 Service Authentication

Individual service authentication—access to privileged accounts on servers, administrative interfaces on switches or routers, etc.—is no different from a remote connection than that required if the individual is attempting a connection from a workstation or server on the LAN. Requiring a login and password is common, but other schemes exist, such as SecureID cards, iLink buttons, etc. All will work over the VPN as if it were on the LAN itself.

For the purposes of this paper, it will be assumed that, within an existing LAN, authentication has already been addressed, at least to the level acceptable for the site.

That leaves authentication of the requestor of a VPN connection as the unique issue to be addressed for remote administration.

3.1.4 Connection Authentication

To authenticate a remote VPN connection, it is necessary to provide *credentials*. This is either an attribute of the connection that establishes it as coming from an approved requestor, or some information provided by the requestor that could only be known by an approved individual or system. Usually, all are used, or at least a combination of two:

- *IP Address*. Only specific IP addresses, or machines within specific subnets, are permitted to even request a connection.
- *Shared keys*. Usually public-key encryption keys that must be cross-authenticated by the VPN software before establishing a connection.
- *Password*. The requesting system must provide a password and, often, a login ID. This is not the same as any password/login for servers or networking equipment on the target site; it's solely for the establishment of the VPN.

Once the requestor is authenticated, the VPN is established.

3.2 SYSTEMS ACCESS

Exactly how an administrator accesses systems once connected depends, largely, on the target system; the most common are discussed in the following sections. Note that if a system or element of interest isn't discussed, that doesn't mean there's no solution for it—this isn't intended to be a comprehensive list.

3.2.1 Servers

3.2.1.1 Linux/Unix

If the VPN connection is to the router/firewall, either a command-line session may be established using SSH[1], or an X-Windows session using a server on the remote workstation. If it's desired to connect to another server from the Linux/Unix system, the same mechanisms can be used if supported, or such tools as pcAnywhere or VNC[2] running on the Linux/Unix server.

If the connection is directly to the server via a tunneled VPN, it will generally be an SSH session, from which an X-Windows session can be started. Again, access to other servers can be established from Unix/Linux-resident tools such as VNC running under X-Windows.

3.2.1.2 Windows

Router/Firewall VPN sessions can run tools such as Windows Terminal Server clients, pcAnywhere, or VNC from the remote client. SSH or VNC can be used to connect to Linux/Unix systems.

If the connection is directly to the server via a tunneled VPN, it will generally be a Terminal Server or pcAnywhere session, from which access to other servers can be established via pcAnywhere, VNC

3.2.1.3 Novell

Remote administration tools such as RCONSOLE or ConsoleOne can work over a VPN link to a router/firewall. If the connection is to a specific system such as a Unix/Linux or Windows, provision must be made to support *port forwarding* to the Novell server(s).

It is generally not a practice to connect from a Novell server to other servers for administrative sessions.

3.2.2 Desktop/Workstations

3.2.2.1 Windows

Windows workstations can be managed via Windows Terminal Server, pcAnywhere or VNC from Windows servers, and via pcAnywhere or VNC from Unix/Linux servers.

3.2.2.2 Unix/Linux

The same tools used for Unix/Linux servers are applicable to Unix/Linux workstations.

[1] Classic telnet can be used, but it is increasingly deprecated for use even within a LAN due to the fact it is absolutely unsecure.

[2] VNC stands for Virtual Network Computer, an Open-Source package originally written and released by Olivetti Research Labs and now maintained by AT&T Laboratories Cambridge. For more information, see http://www.uk.research.att.com/vnc/.

3.2.2.3 Macintosh

Macintosh, both pre- and post-OS X can be accessed and access other systems using tools such as Timbuktu Pro.

3.2.3 Network Components

Most network components accept either telnet or web (HTTP) connections. In either case, these can be effected either directly through firewall/router VPN links, or from the target server for a tunneled VPN link.

3.3 AUDITING

Generally, firewall/routers can be programmed to log VPN sessions, as can Unix/Linux and Windows servers accepting tunneled VPN sessions. Rules generating log entries for remote sessions should be established prior to putting a system in production; moreover, logs should be sent to a printer or a write-only device to guarantee they are tamper-proof.

3.4 REMOTE NOTIFICATION

An adjunct to remote administration is provision for remote notification in case of system events that require the attention of an administrator. Actually, this is a useful feature for on-site support as well.

Remote notification can be triggered by a number of utilities, including Simple Network Management Protocol (SNMP) clients, Tivoli, or custom log file auditing scripts. In all cases, the usual response is to cause the sending of an E-mail or, more usefully, a text message or page to a beeper, or a text message or call to a telephone or wireless number.

4. CONCLUSION

The capability to carry out remote administration is a valuable tool for system and network administration. Tools and techniques for doing this in a reliable and secure manner are readily available for little or no cost, and add greatly to the range of support options available to the site administrators.